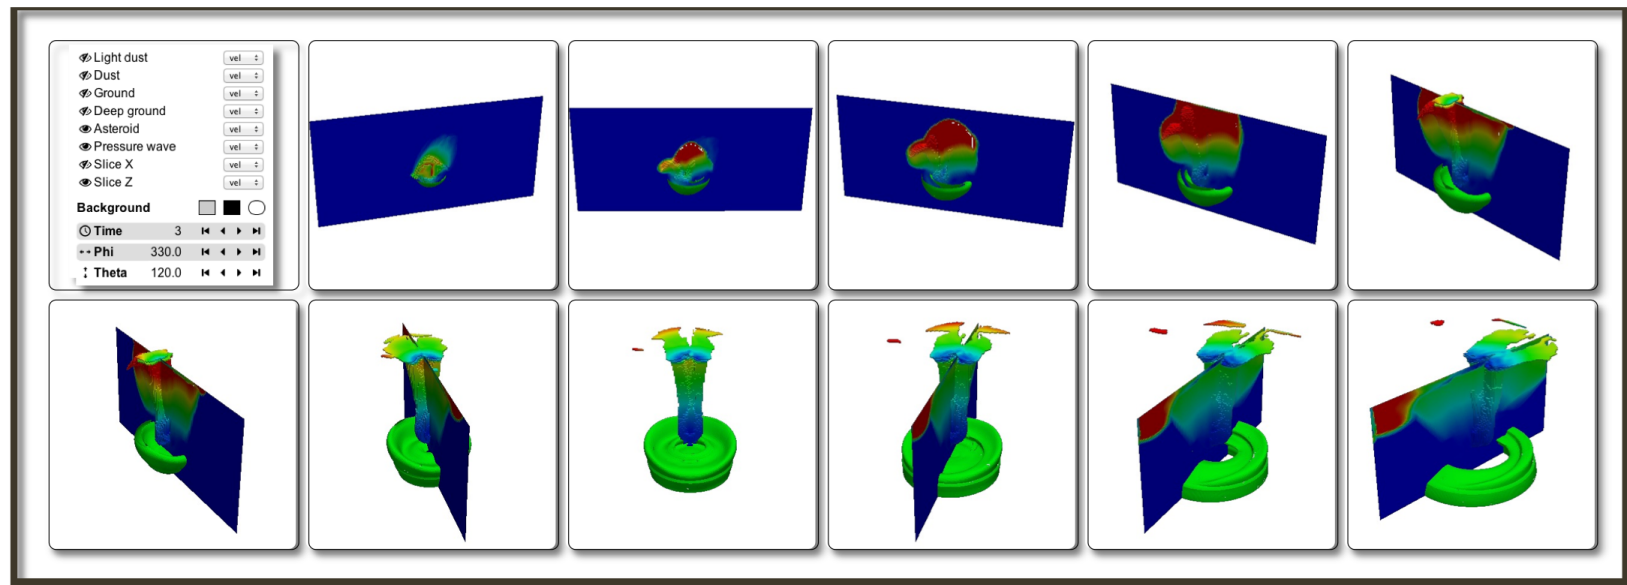


# Implications of Numerical and Data Intensive Technology Trends on Scientific Visualization and Analysis



# Technology Trends in Numerically and Data Intensive Computing

## **Numerically intensive Trends**

Hardware: Exascale challenges and solutions

## **Data intensive Trends**

Software: Cloud challenges and solutions

# Trends for HPC Scientific Visualization and Analysis

Relentless increase in data sizes  
3 orders of magnitude every  
ten years

Adapting to changing  
infrastructure

Shared memory, clusters,  
threading, cloud



Advancing the fundamentals

Improved end-to-end workflow  
and cognitive understanding  
How about the user  
experience?



# Responding to the Trends: ParaView

- An open-source, scalable, multi-platform visualization application
- Support for distributed computation models to process large data sets
  - Billions of AMR cells, Scaling test over 1 Trillion cells
- Used by academic, government and commercial institutions worldwide
  - Downloaded ~100K times per year
  - Developed by Kitware, LANL, SNL...
- Originally designed to support a post processing workflow
  - Simulations save data to storage and scientist interactive visualizes results



<http://paraview.org>



# Numerically Intensive Trends: Exascale Computing – The Vision

Achieve order  $10^{18}$  operations per second and order  $10^{18}$  bytes of storage

Address the next generation of scientific, engineering, and large-data problems

1,000X capabilities of today's computers with a similar size and power footprint

Set the US on a new trajectory of progress – towards a broad spectrum of computing capabilities over the next decade

## Productive system

- Usable by a wide variety of scientists and engineers
- “Easier” to develop software & management of the system


## Based on marketable technology

- Not a “one off” system - Scalable, sustainable technology
- Deployed in early 2020s



# Potential Exascale System Architecture

With a cap of \$200 M and 20 MW

Feature	2013 Titan Computer	2023	Difference 2013 & 2023
System Peak	27 Pflops/s	1 Eflop/s	O(100)
 Power	8.3 MW	20MW	2.5x
System Memory	0.7 PB	64 PB	O(100)
Node Performance	1.5 TF/s	15 TF/s	O(10)
Node Memory BW	0.2 TB/s	4 TB/s	O(10)
Interconnect BW	0.008 TB/s	0.4TB/s	O(100)
Number of Nodes	18688	100000	O(10)
Total concurrency	50M	O(billion)	O(100)

Power is very costly: 1 MW = ~ Million dollars  
Without intervention on track to 200MW for Exascale

# Data Access Delay

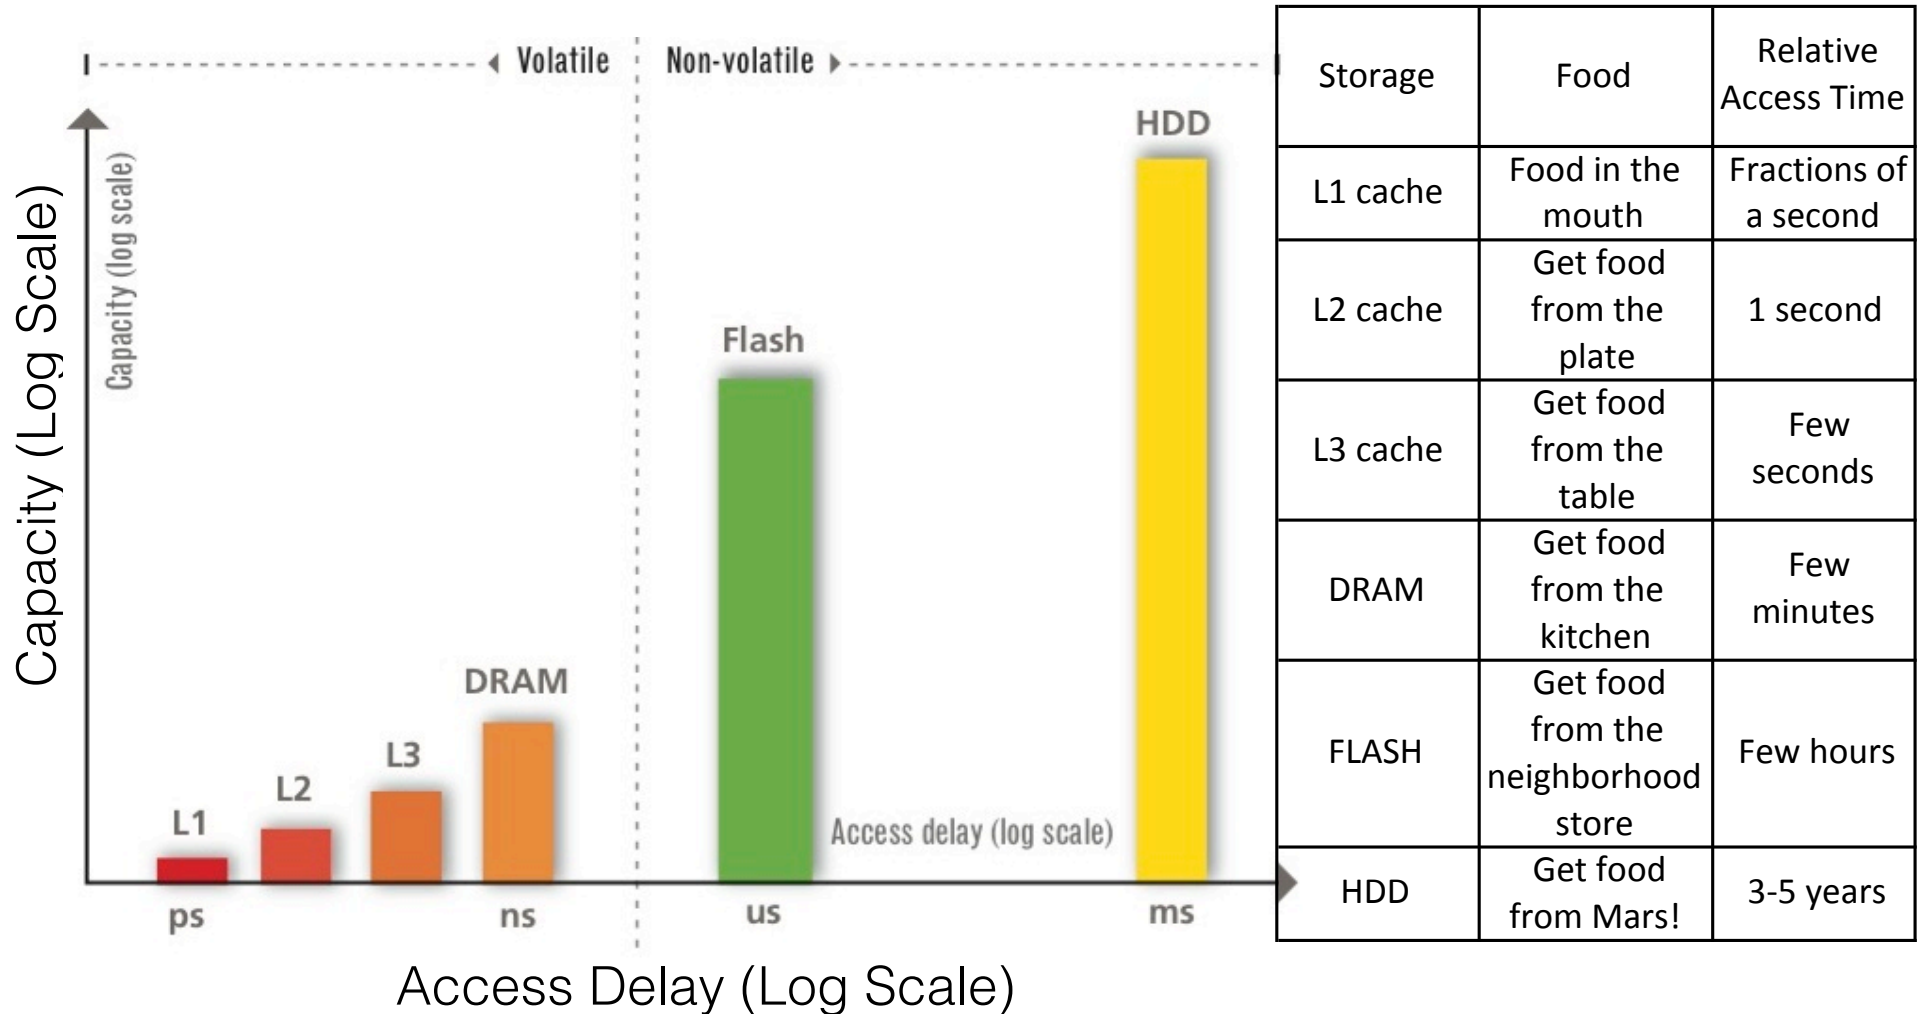


Diagram and Table from “Taming the Power Hungry Data Center”, Fusion I/O.

Implication: The traditional post-processing approach is becoming unworkable at extreme scale

- Temporal simulation snapshots are saved at longer intervals
  - Full checkpoints are costly - less temporal data available for analysis
- Rate of improvement of rotating storage is not keeping pace with compute
  - Power, cost and reliability are becoming significant issues

# Implication: Transition to an in situ focused approach

- In situ saves reduced-sized data products during simulation run
  - Benefits:
    - Save disk space
    - Save time in post-processing analysis
    - Produce higher fidelity results
- Automatic visualization and analysis during the simulation run
  - Prioritized by scientist's importance metrics
- Identify specific analysis questions
- Help manage cognitive and storage resource budget




## Implication: Significant in situ data reduction

Algorithm	Reduction
Data parallelism	Handle large datasets Make reduction possible
Multi-resolution	Make focused exploration possible
Visualization and analysis operators (isosurface)	A dimension reduction
Statistical sampling	1-2 orders of magnitude
Compression	1 order of magnitude
Feature extraction	2 orders of magnitude

# Sampling


- Random sampling provides a data representation that is unbiased for statistical estimators, e.g., mean and others
- Since the sampling algorithm is in situ: accuracy metric(simulation data, sampled representation)

## Plotting




Red is 0.19% sample data, black is original simulation data

## Feature Extraction: Halo Finding




- . The red, green, and blue curves are 0.19%, 1.6%, and 12.5% samples. . The black curve is the original data.
- . Calculate the halo mass function for different sample sizes of  $256^3$  particles

## Example: Visual Downsampling





Cosmology visualization in ParaView



# In Situ Compression with Quantified Accuracy


- *In situ* compression of simulation data
  - Use JPEG 2000 to compress data
  - Quantify the maximum/L-infinity norm) data quality for scientific analysis
- Measure the maximum point error
  - Guarantee accuracy to x decimal places
  - Accuracy Metric (Simulation data – Compressed representation)
- User can trade read I/O time vs. data accuracy in a quantifiable manner






# Isovalues on Compressed Simulation Data with Bounding Error - (32 bits, 3200x2400x42, 1.4 GB)


0.25 bits  
10.8 MB




0.5 bits  
21.6 MB



1.0 bits  
43.3 MB




2.0 bits  
86.5 MB



# Implication: Automated Algorithms

## Adaptive focus based on selected scientific metrics




- Create adaptive analysis-based grid
  - Histogram at each grid element
    - Across all axes (spatial, value, multivariate)
- Use for spatial, temporal selection
  - Cameras, storage, feature identification


xRage 1209.01

ito296

setup by Bob Weaver



.




Time: 0.000000 s


06/01/2013 10:50 AM

# Sampling Using Analysis Driven Refinement (ADR)

- Recursive metric-based refinement
- Multidimensional



Sampling in Time



Sampling in Space

# Data Intensive Trends: Cloud Computing

## The NIST Definition

- A model for enabling ubiquitous, convenient, on-demand network access to:
  - a shared pool of configurable computing parallel resources
    - (e.g., networks, servers, storage, applications, and services)
  - rapidly provisioned and released with minimal interaction
- <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>



# The NIST Definition of Cloud Computing

## Essential Characteristics

---


- On-demand self-service
- Resource pooling / Multi-tenancy (multiple jobs)
  - Virtualization
- Rapid elasticity
  - Scale rapidly commensurate with demand
- Measured service / Cost model
  - Resource usage is automatically monitored, controlled, and reported, providing transparency

# The NIST Definition of Cloud Computing

## Essential Characteristics

---

- Levels of cloud service
  - Infrastructure
  - Application
- Private cloud is an option...



Axis	Sub-axis	Numerically Intensive	Data Intensive
Hardware	Nodes and Interconnect	High performance and power	Lower performance and power
	Storage	Separate, independent	Integrated
SW	Synchronization	Tightly coupled	Loosely coupled
	Reliability	Checkpoint restart	Replication
Workload	Number of Users	<u>Single per node</u>	<u>Multiple per node</u>
	Data	Dynamic, heterogeneous (unstructured grid)	Static, homogeneous (text, images)
	Algorithms	Global	Distributed
	User Interface	<u>Complex Application</u>	<u>Simple Web</u>
	Data Model	<u>Files</u>	<u>Database</u>
Workflow	Scheduling	Batch	Interactive
	Analysis	Offline post-processing	Online
	I/O	Bulk parallel writes	Streaming writes

# Implications of Cloud Computing on HPC Visualization and Analysis

Multi-billion dollar market

- Leverage, collaborate and support

Virtual machine (VM) encapsulates a simulation with defined inputs/outputs

- Cloud infrastructure services require VM
  - Provenance - full lineage of data/process/environment
  - Resilience – follows from provenance
  - Data compression – VM and input deck instead of data
  - To do: Reduced VM size and VM composition

# Implications of Cloud Computing on HPC Visualization and Analysis

## Data-oriented applications As an approach to massive data

- Beyond Map-Reduce
  - Environments – Spark
  - Scalable databases – Impala, MongoDB
  - Data analytics products

## User/task-centric applications

- Cloud enables mobile/web
- Focus on usability and simplicity



# Inspiration: Image Database Approach Cinema

## Challenge

In situ is a batch process

Concern that exploratory aspect of analysis will be lost

## Idea

Store *many* images that sample the visualization parameter space

In less than the space needed for a single scientific data dump

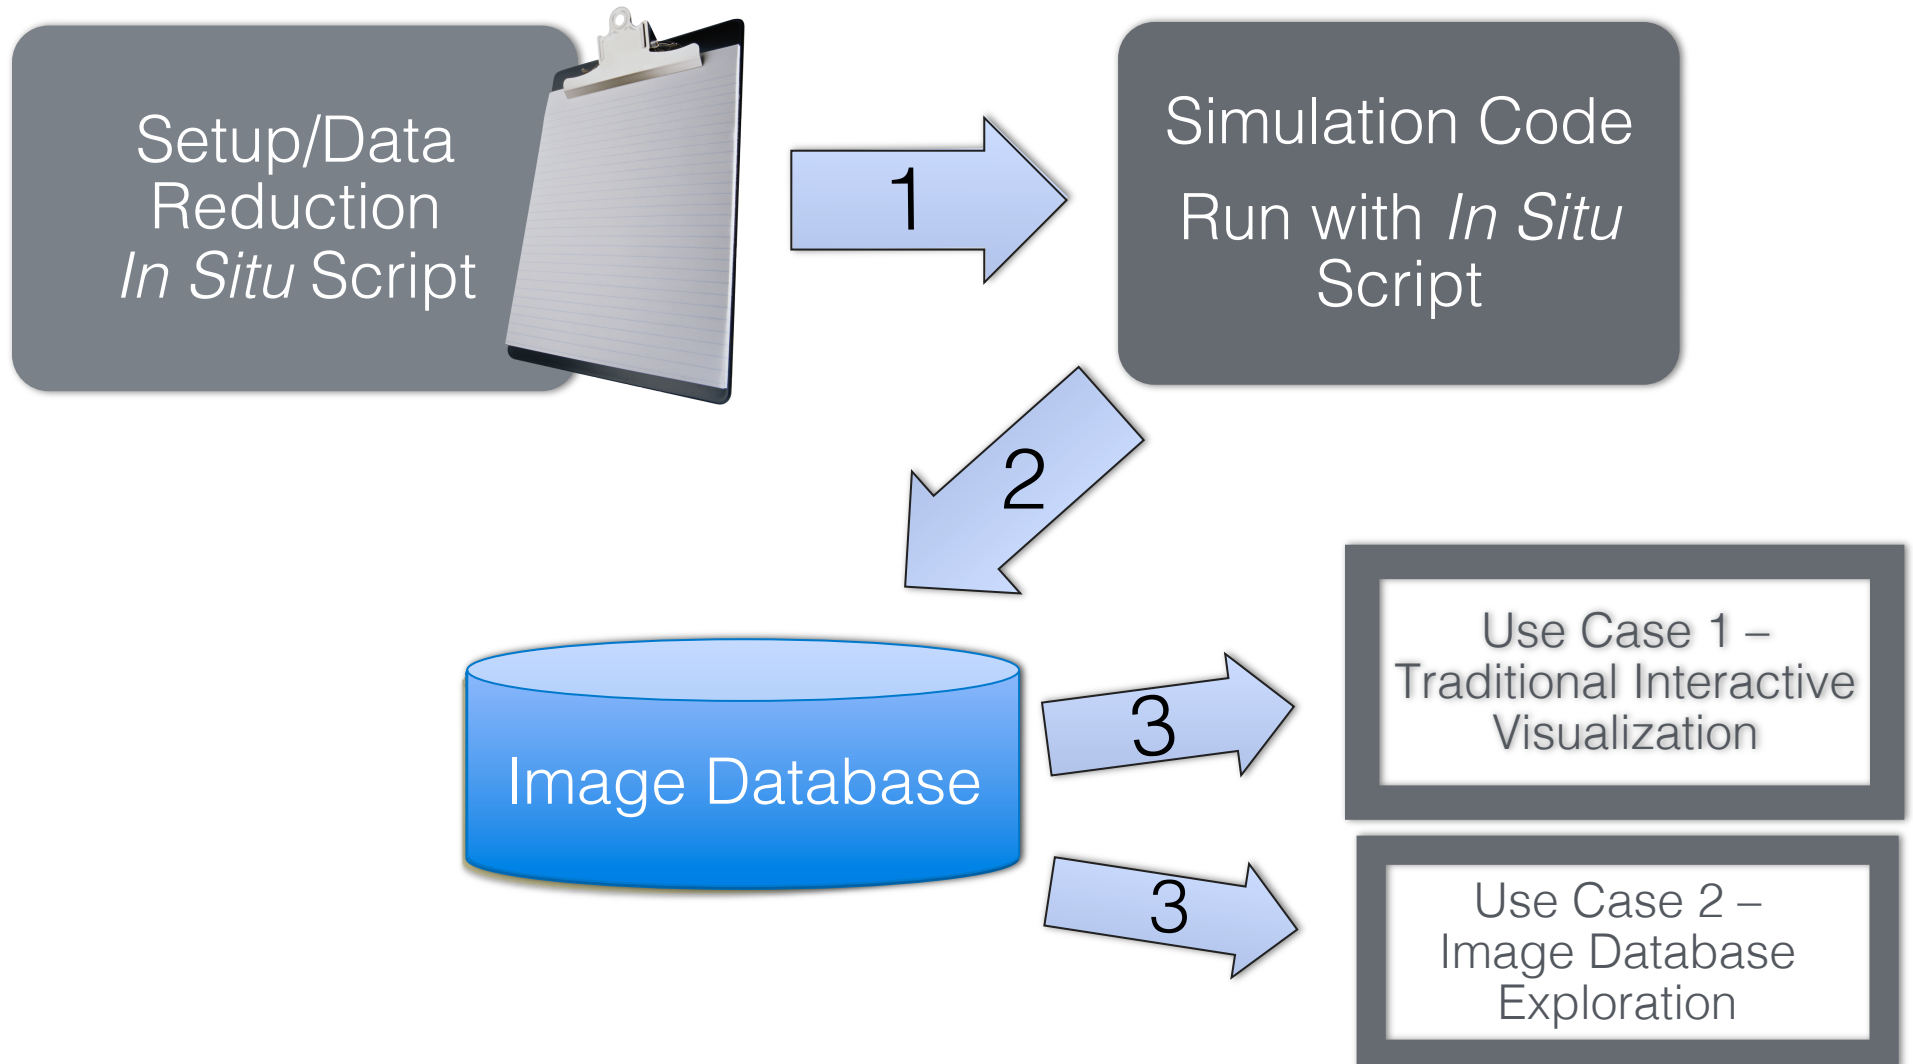
Ex: Cameras, operations, parameters

Create an image database from in situ analysis  
Post-processing exploration of image database

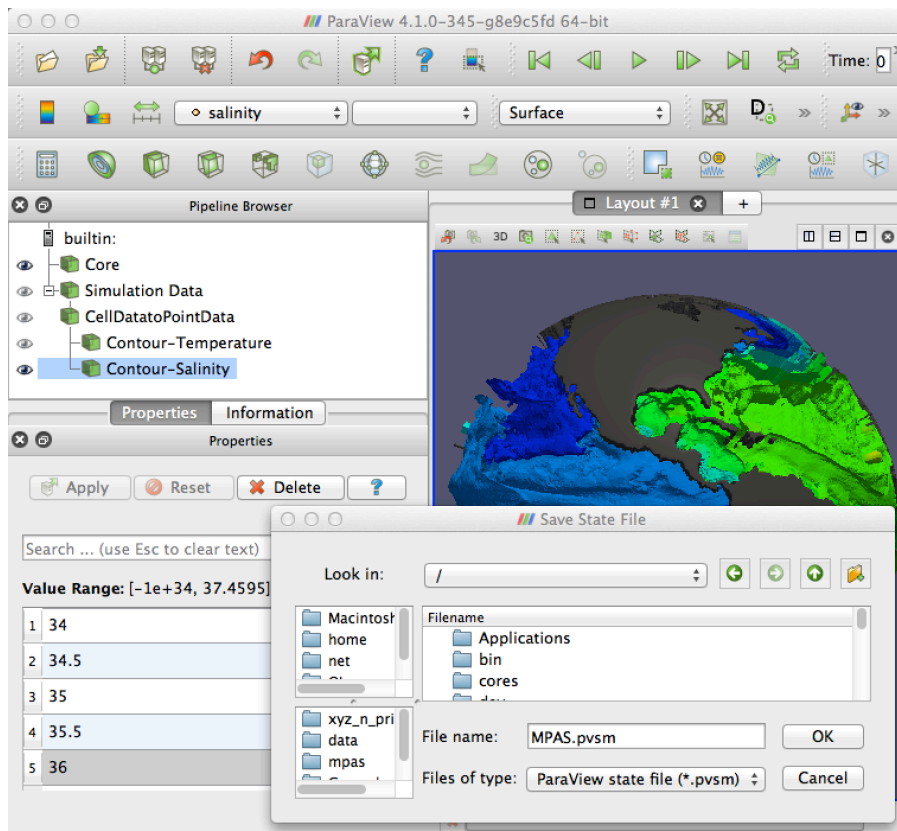



Mega	Giga	Tera	Peta	Exa
$10^6$	$10^9$	$10^{12}$	$10^{15}$	$10^{18}$
Image speed	Storage & network speed	Operations speed	Operations speed	Operations speed

# Cinema Workflow



# Setup /Data Reduction Phase




- Interactively create or reuse a visualization pipeline
  - Contains all operations
  - Specifies information needed to generate images for the database

# Setup / Data Reduction Phase

**Upload visualization pipeline state**  MPAS.pvsm

Pipeline

- **Earth core**
  - Color by ☒ 0.5, 0.5, 0.5
- **Simulation data**
  - Simulation parameters
    - Simulation timesteps
    - Output frequency
- **CellDataToPointData**
  - **Contour**
    - Parameters
      - Contour by
      - Contour values
    - Color by
      - ☒ Temperature ☒ Salinity ☐ Density
      - ☐ Pressure ☐ 0.5, 0.5, 0.5
  - **Contour**
    - Parameters
      - Contour by
      - Contour values
    - Color by
      - ☒ Temperature ☒ Salinity ☐ Density
      - ☐ Pressure ☐ 0.5, 0.5, 0.5



Set camera and operator parameters to visualize

# Image Database

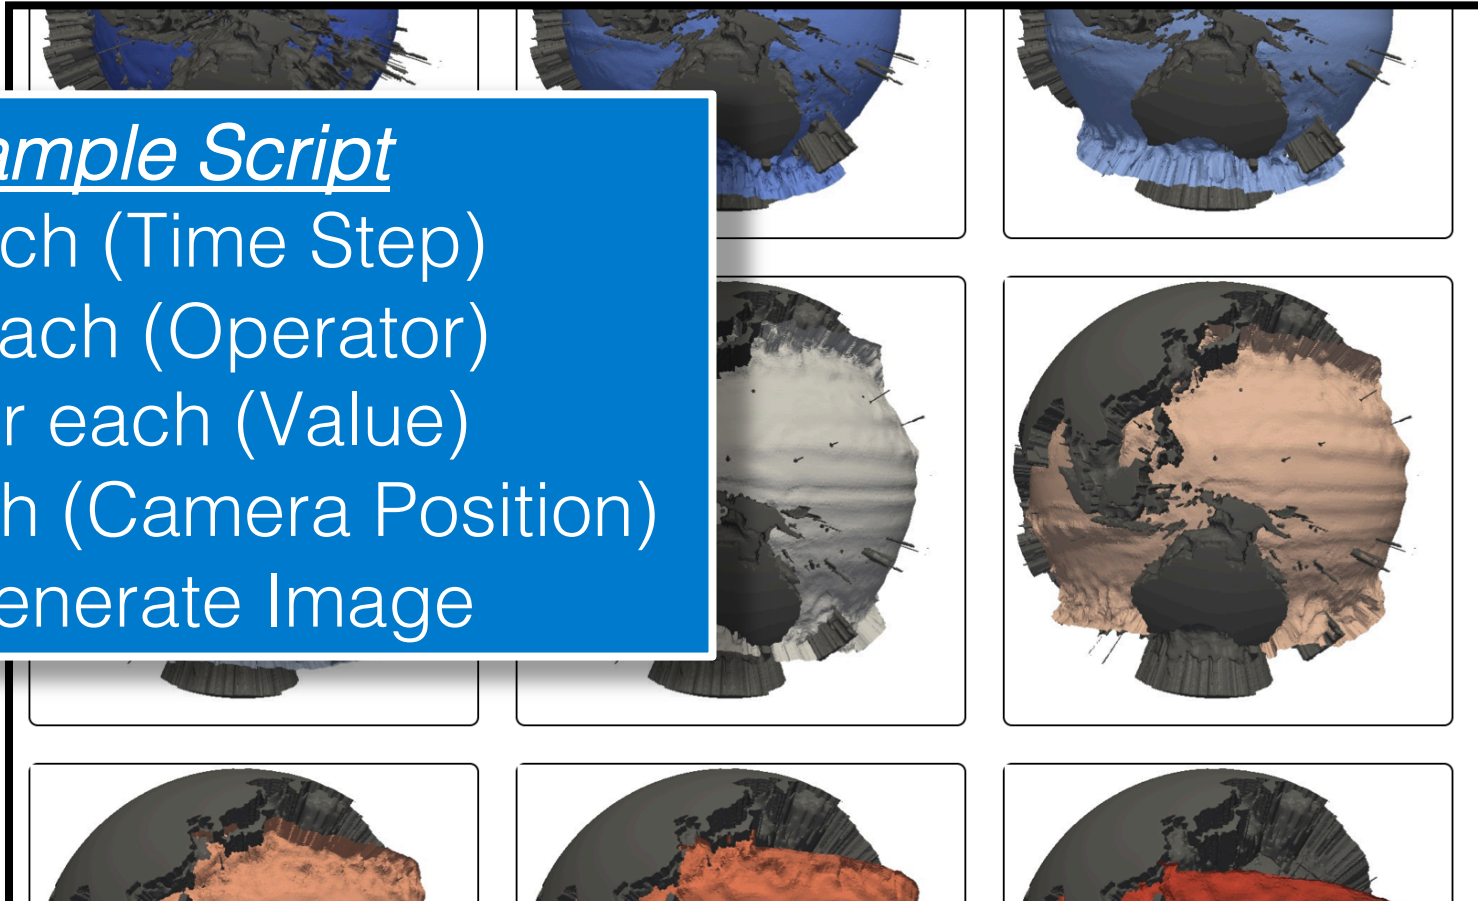
Simulation Code  
Run with *In Situ*  
Script

2

Image Database

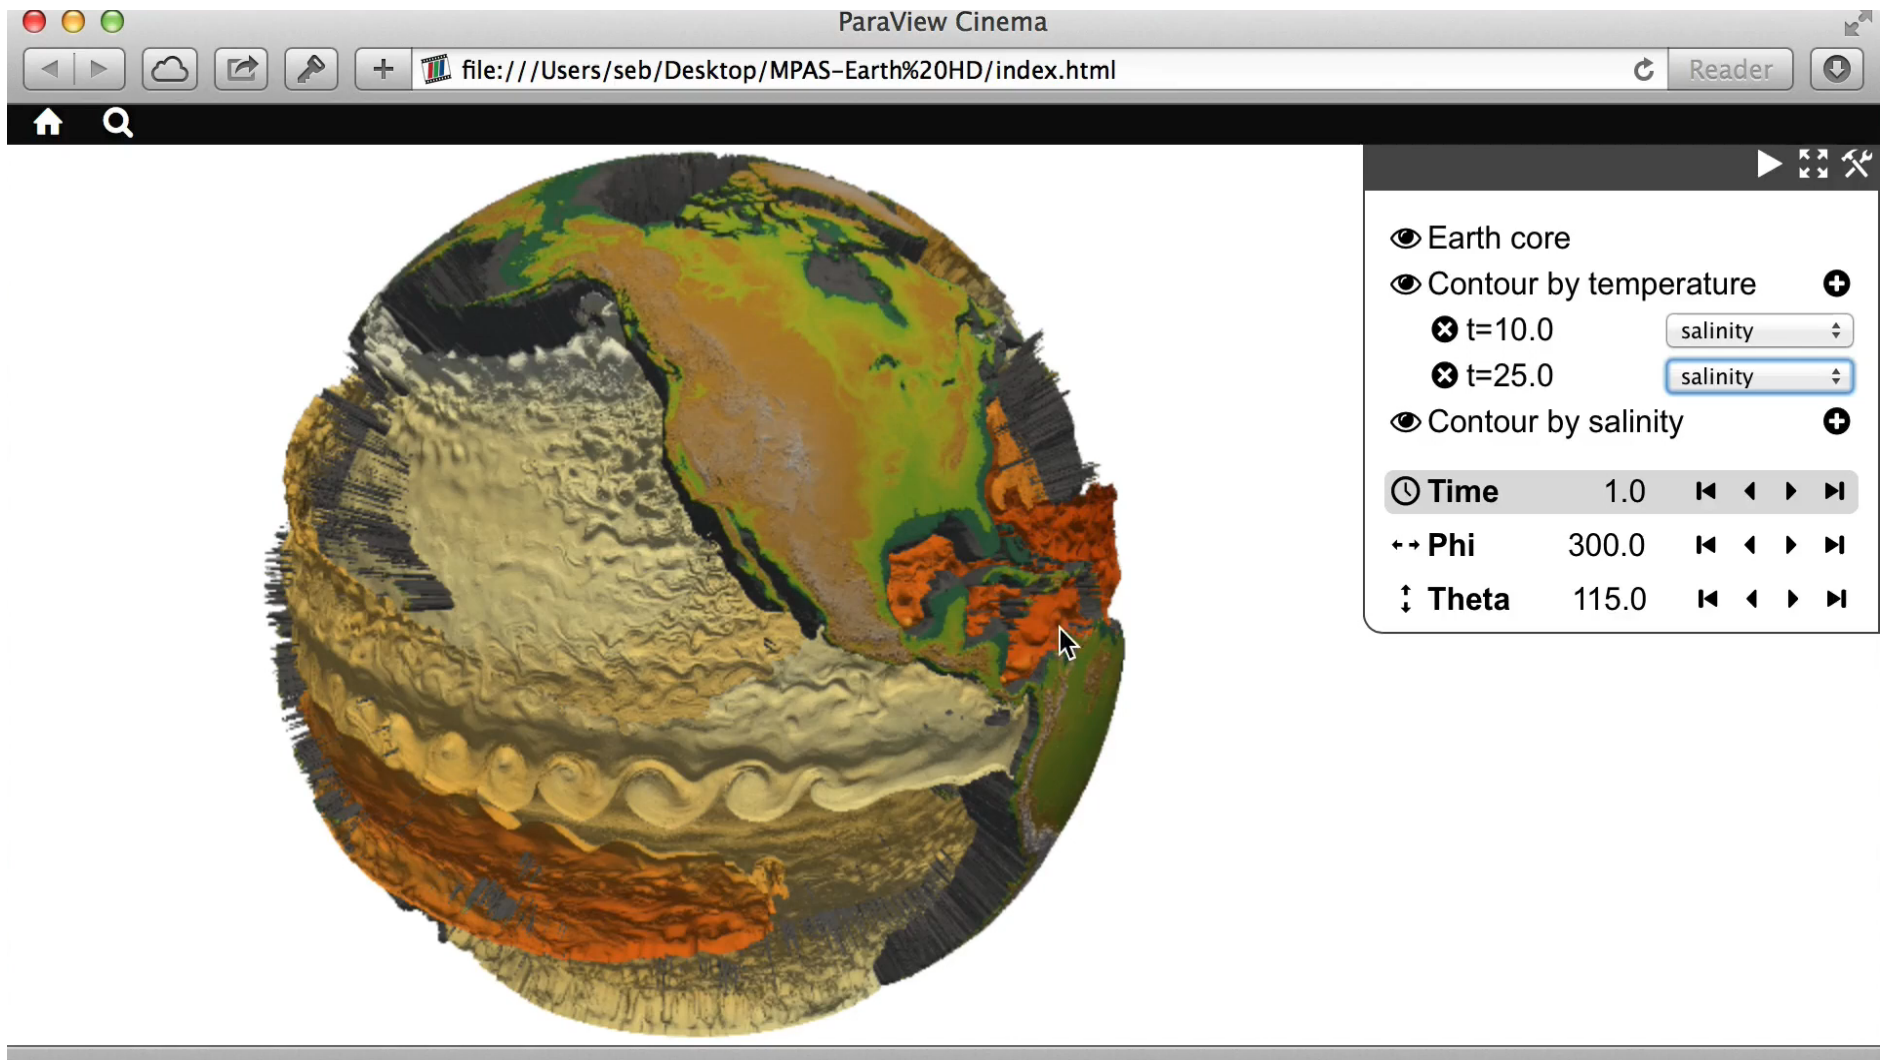
## Example Script

```
For each (Time Step)
  For each (Operator)
    For each (Value)
      For each (Camera Position)
        Generate Image
```





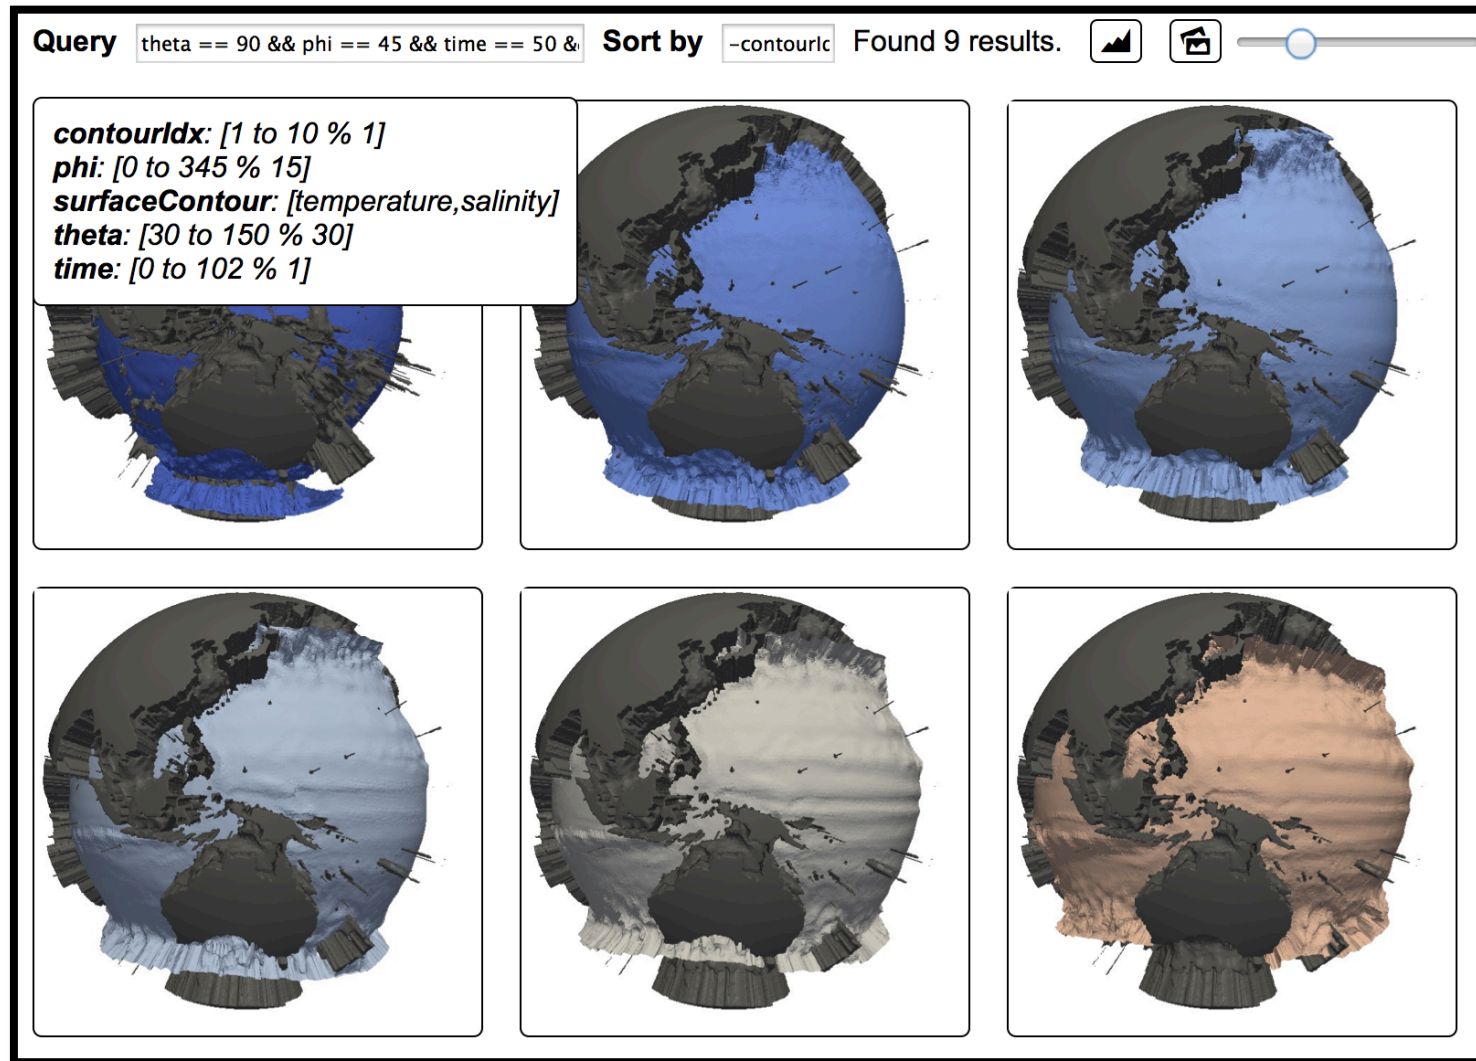
# Use Case 1 – Traditional interactive exploration



In all videos in this presentation:

Processing, combining and showing images from the image database  
No raw scientific data is read, no geometry is created during viewing

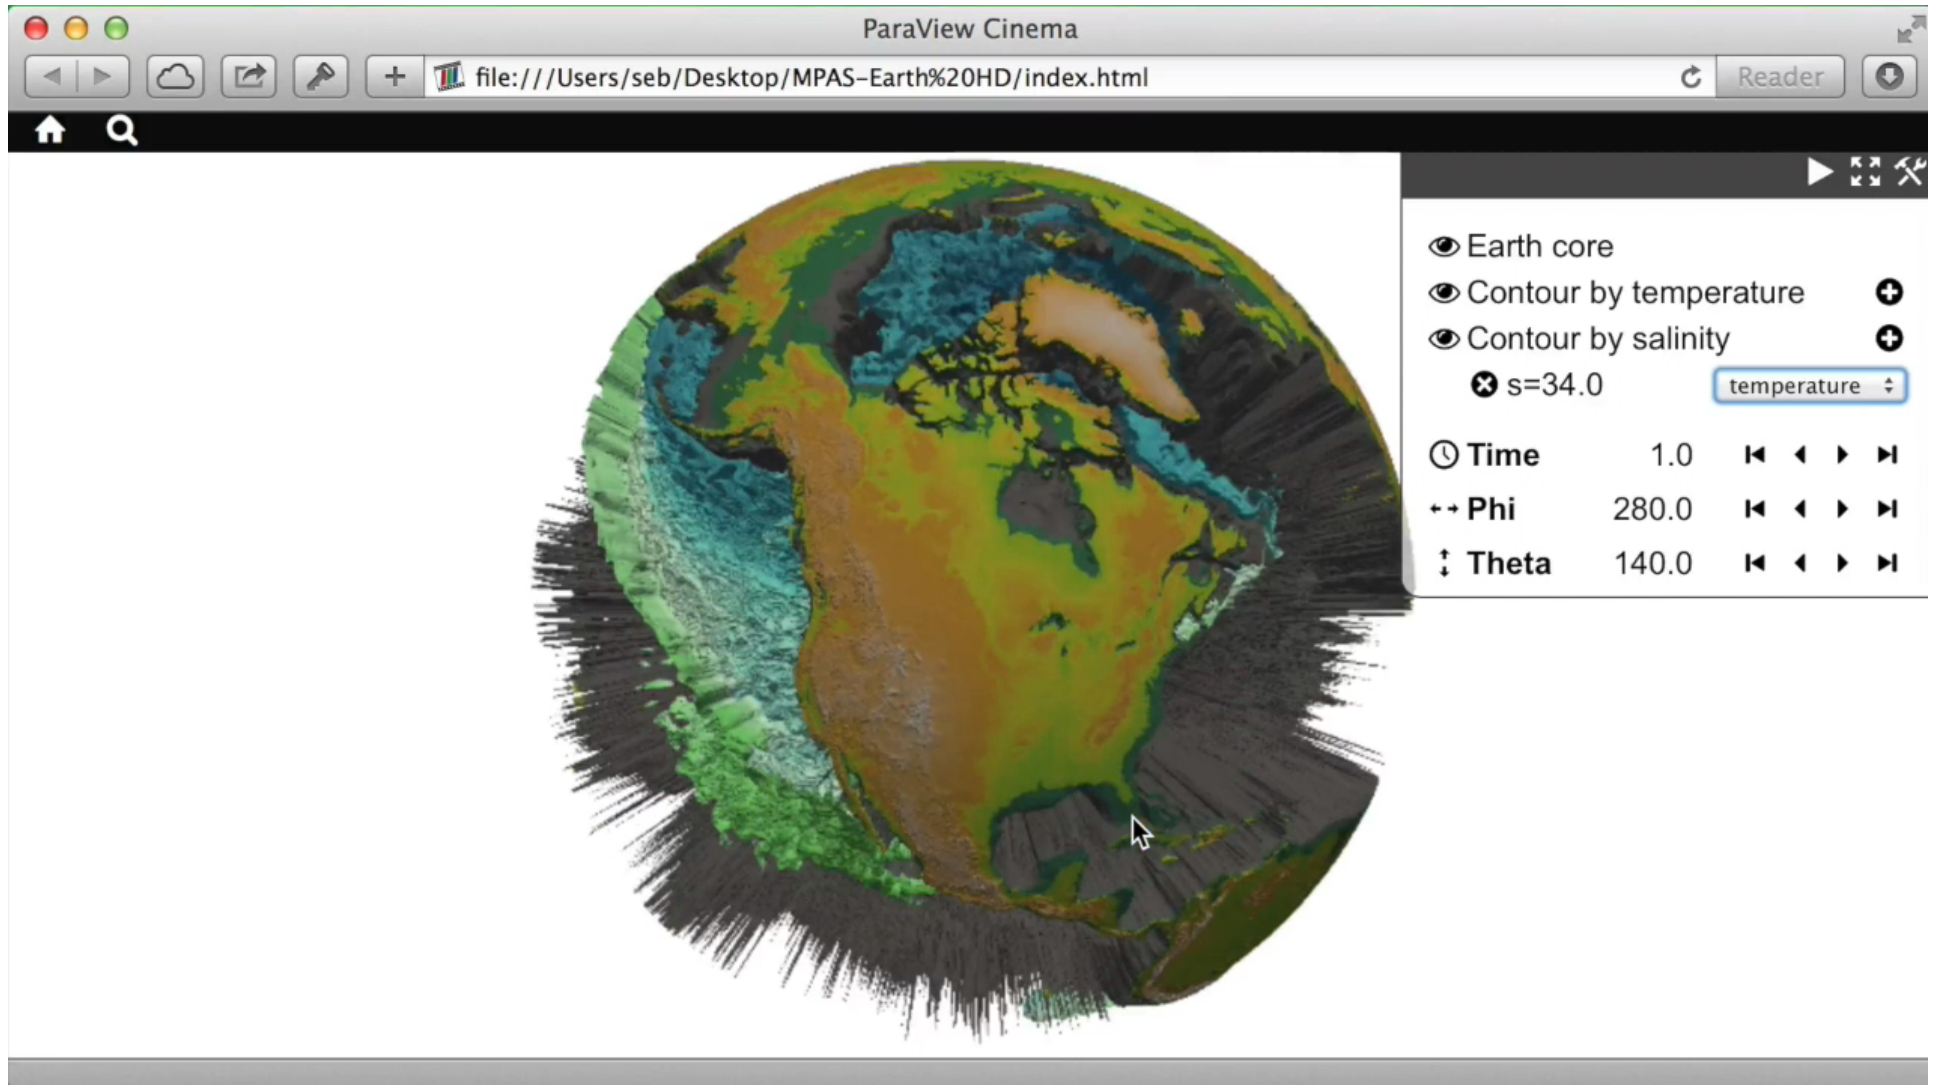
## Use Case 2 - Image database exploration



Traditional key-value pair queries  
Keys: Camera (phi, theta), time, operator parameters




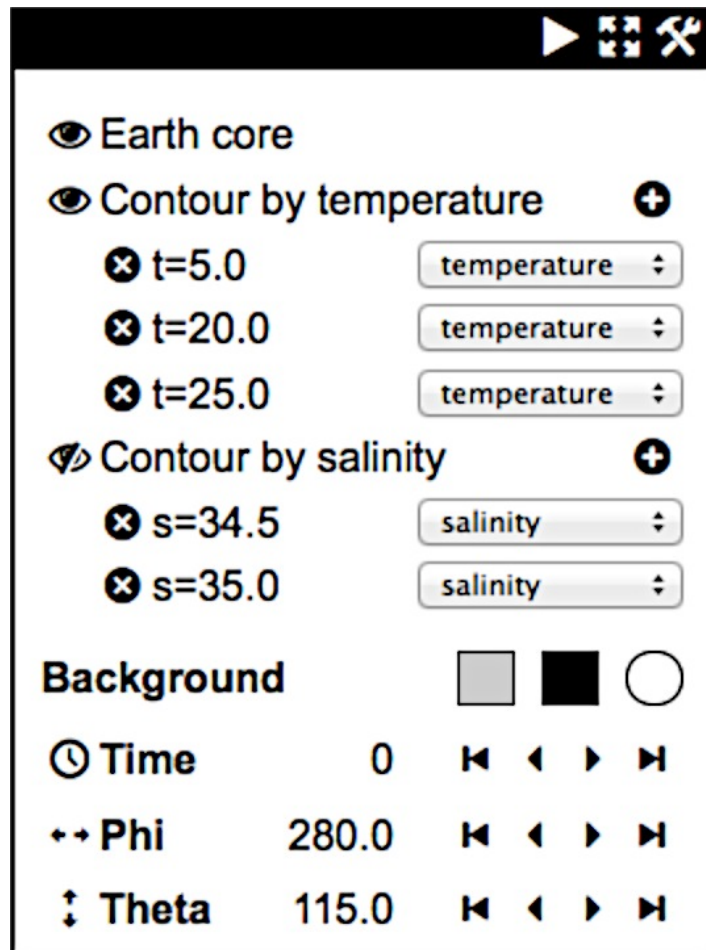
# Use Case 2 – Image database exploration



# Image-based approach reduces analysis exploration bias

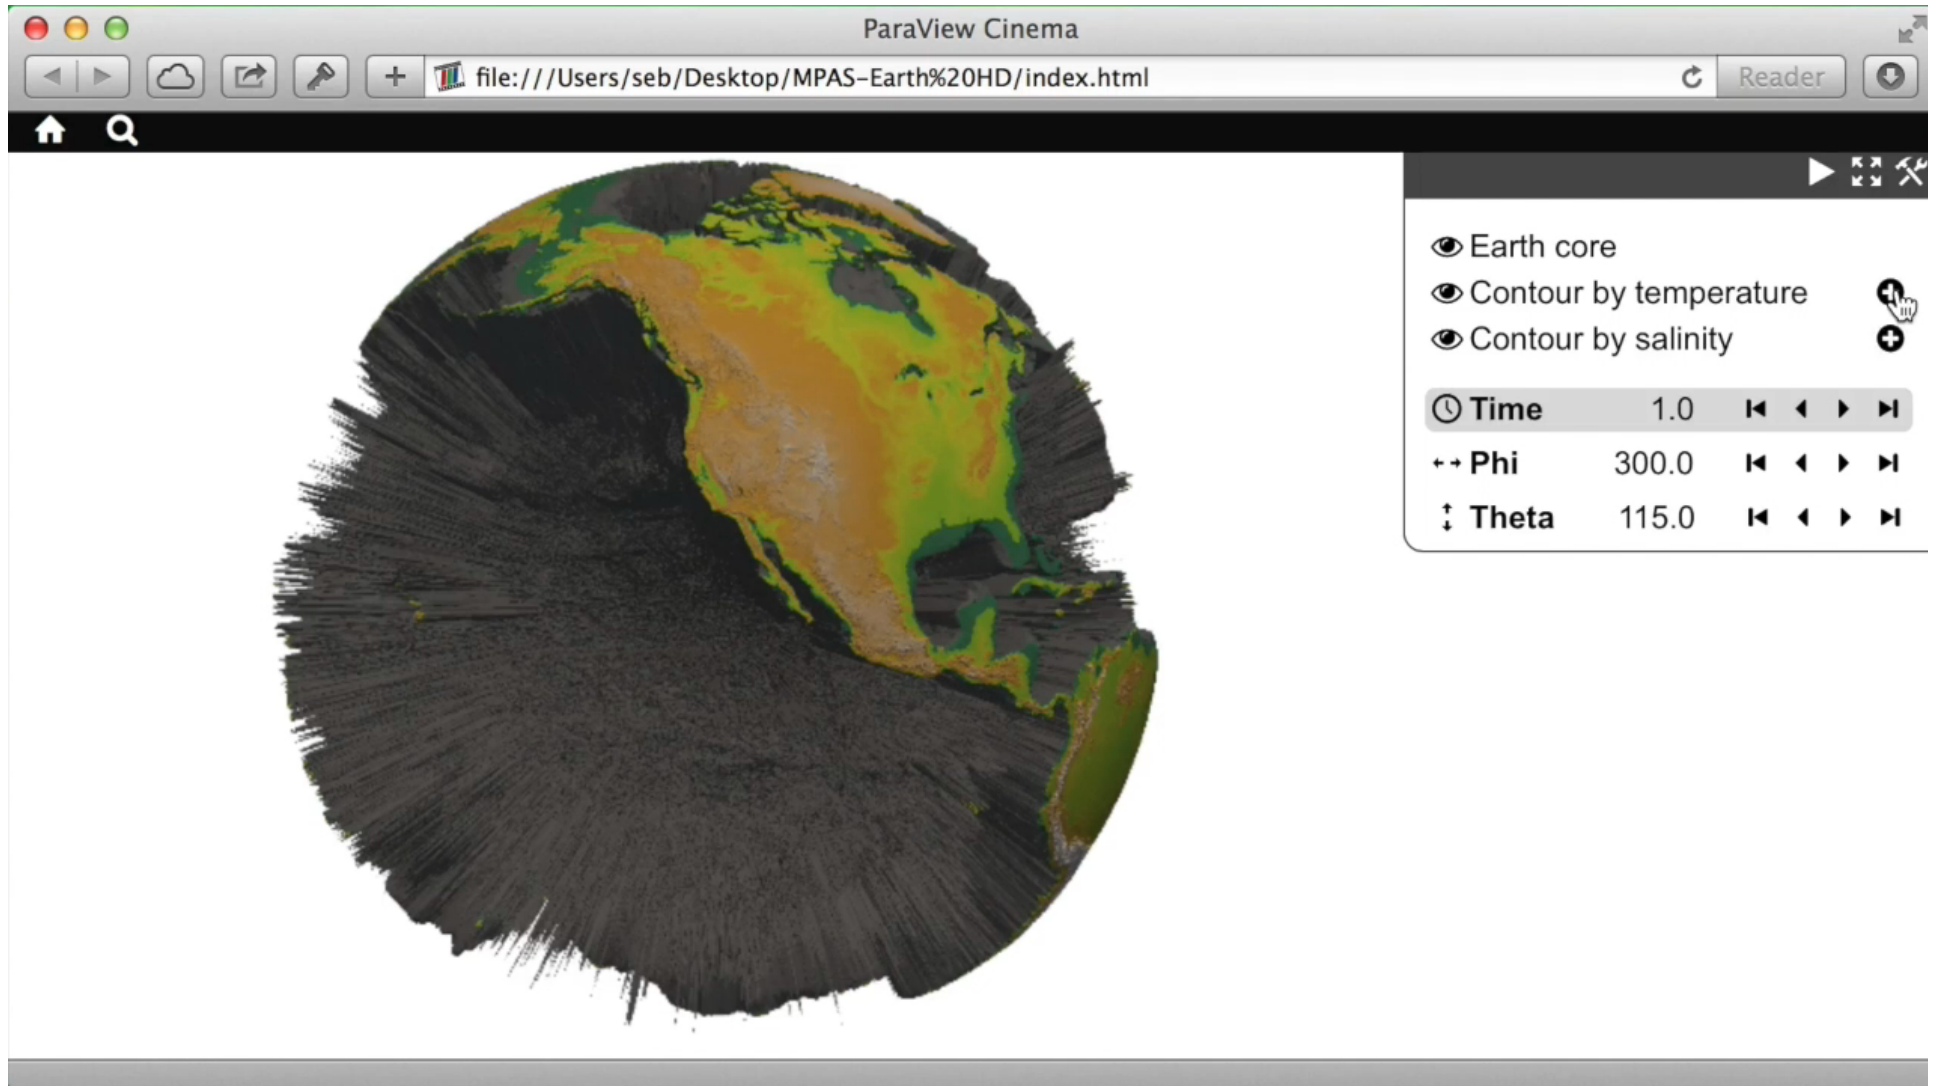
- Traditional post-processing approach
  - Generate visualization and analysis result upon user request
  - User wait time is extremely variable
    - Rendering (seconds)
    - File system accesses (minutes)
  - Creates inherent bias in what is explored
    - For example: little significant interactive temporal analysis
- For an image-based approach
  - All “operations” take the same amount of time
    - Reduces bias of what get explored

# Use Case 3 – Creation of new visualizations



- Use real time image compositing to build new pipelines
  - Image representation: Color & depth buffer
- Multitude of combinations/visualizations possible



# Use Case 3 – Creation of new visualizations



- Scientists can quickly create “arbitrary” pipelines to answer their analysis questions



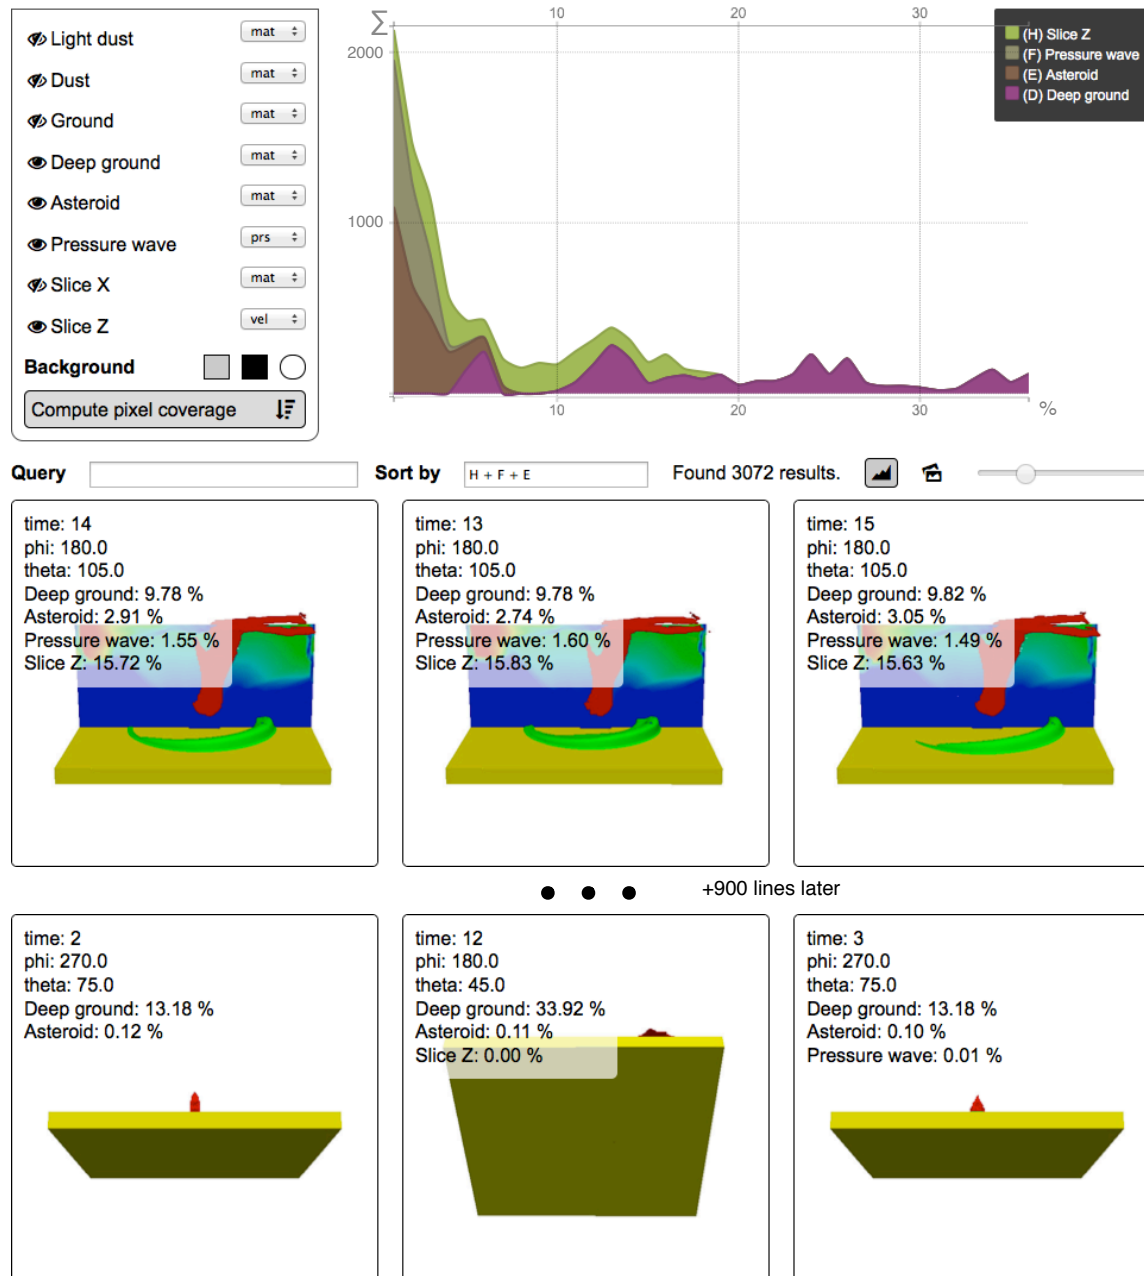
## Use Case 2 & 3 – Content based image search




Percent of image covered

- What image in the database contains the “best” view of a collection of visualization objects?
  - Each image/pixel contains a list of the order/visibility of the objects for each view
  - Pixel coverage is calculate for all views and objects

# Use Case 2 & 3 – Content-based image search






- Databases
  - Plasma Code /Intel Ray tracer, MPAS/Cinema in-situ, HACC Cosmology data
- Code examples
  - Coupled MPAS/Cinema to create new databases

<http://datascience.lanl.gov/Cinema.html>



# Weak Scaling of XRage Simulation and In Situ Analysis




In situ analysis of 10 contour objects and background  
Image size of 500x500

Summary: Scalable in situ performance to generate database

# Disk usage reduction

## Full XRage data files versus in situ




Summary: Orders of magnitude data saving with Cinema approach

# Conclusions

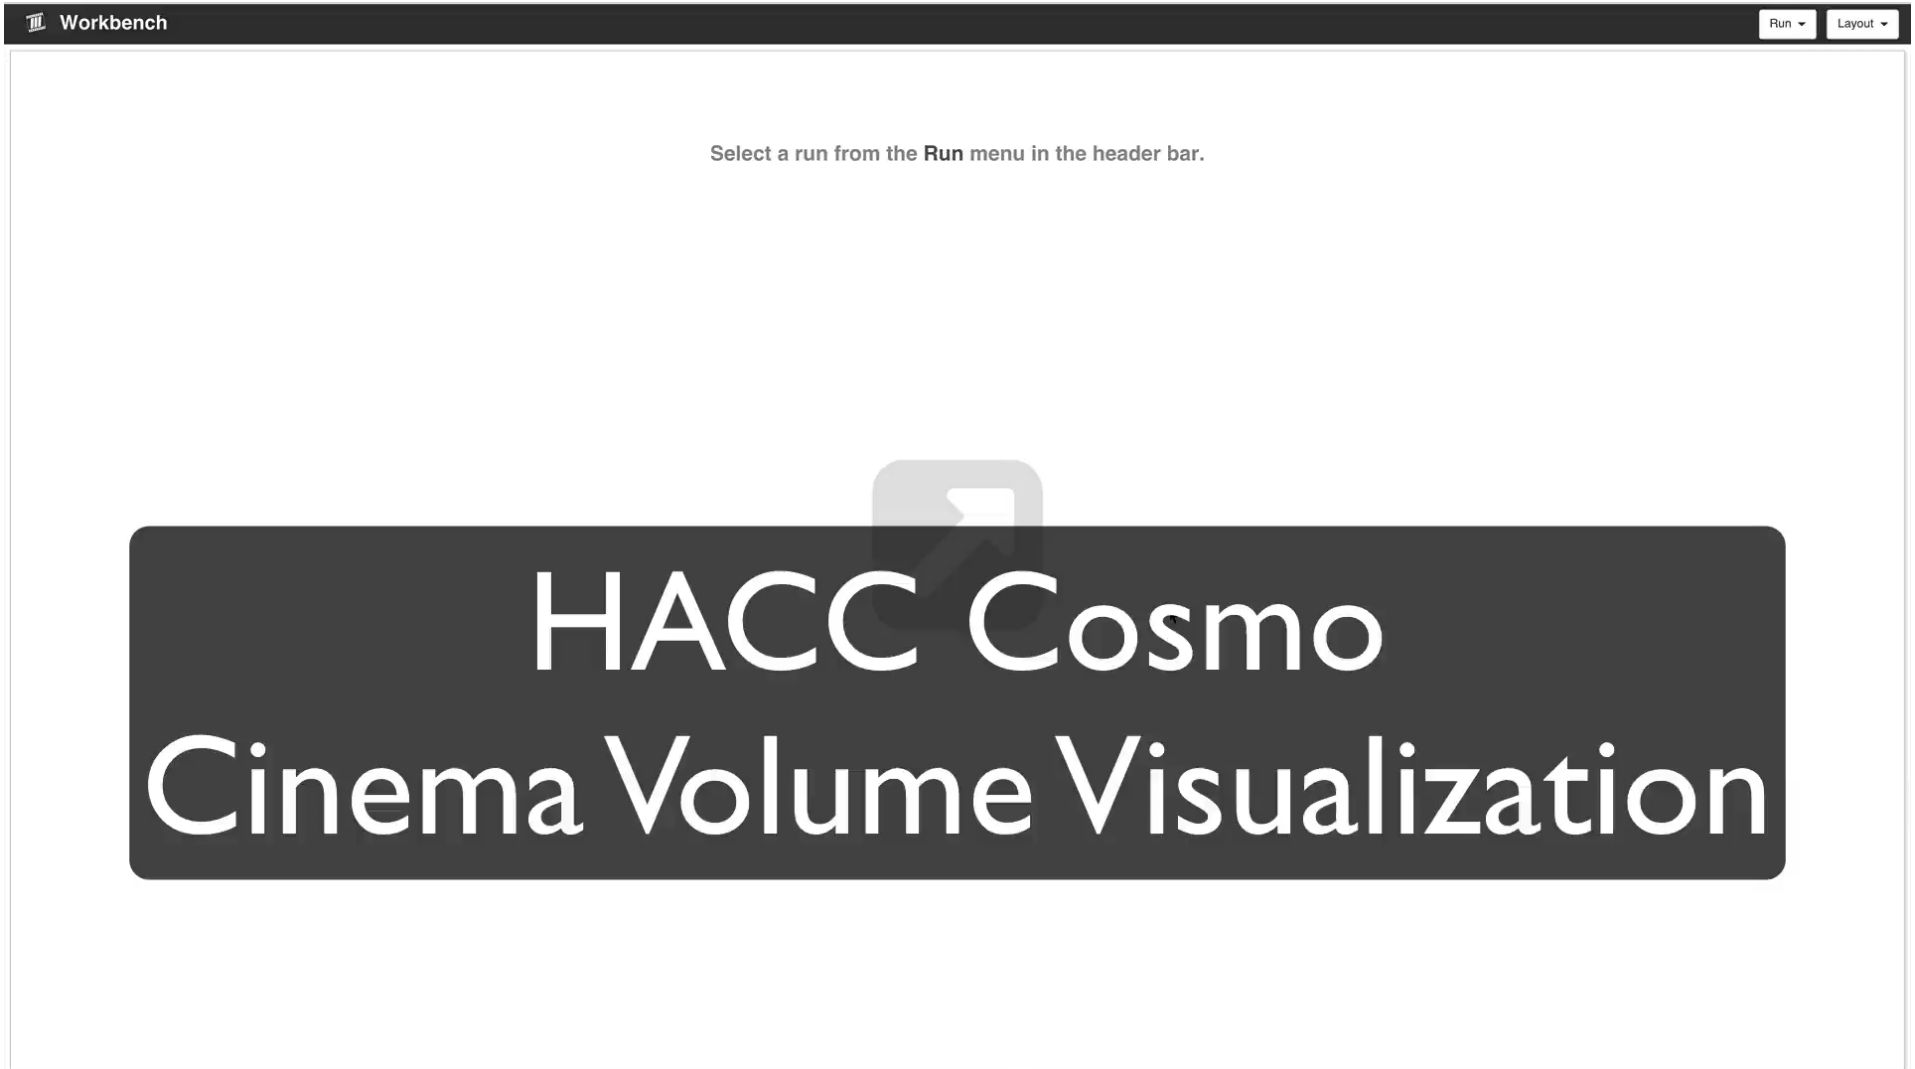
- Next steps: <http://datascience.lanl.gov>
- In situ workflows are required for exascale
  - Benefits over traditional post-processing approach
  - Sampling is key
- Reduced simulation data approach
  - Error quantification is possible
- Image database approach
  - Offering unique interactive exploration options
    - Database search

Acknowledgements - This work was funded by Dr. Lucy Nowell, Program Manager for the Advanced Scientific Computing Research (ASCR) program office in the Department of Energy's (DOE) Office of Science and Thuc Hoang, Program Manager, DOE NNSA.



Curtis Canada, Patricia Fasel, Li-Ta (Ollie) Lo, John Patchett, David Rogers, Francesca Samsel, Christopher Sewell, Jonathan Woodring, Garrett Aldrich, Ayan Biswas, Chris Bryan, Lalindra De Silva, Daniel Hill, Sidharth Kumar, Evgeni Makevnin, Dennis Mosbach, Jesus Pulido, Andre Schmeisser, Wathsala Widanagamaachchi, Max Zeyen, Connor Dolan, Mike Jacobi, Kalin Kanov, Sidharth Kumar, Peter Ortegel, Kien Pham, Jesus Pulido, Yu Su, Will Vining, Berk Geveci, Patrick O'Leary, Dave DeMarle, Sebastian Jourdain, Greg Abram

# Questions



# Publications

- B. Nouanesengsy, J. Woodring, K. Myers, J. Patchett, and J. Ahrens, “ADR Visualization: A Generalized Framework for Ranking Large-Scale Scientific Data using Analysis-Driven Refinement”, LDAV 2014, November 2014, Paris, France.
- K. Myers, E. Lawrence, M. Fugate, J. Woodring, J. Wendelberger, and J. Ahrens, “An In Situ Approach for Approximating Complex Computer Simulations and Identifying Important Time Steps”, in submission, arXiv: 1409.0909.
- A. Biswas, S. Dutta, H.-W. Shen, J. Woodring. “An Information-Aware Framework for Exploring Multivariate Data Sets.” IEEE Visualization 2013, Atlanta, GA, November, 2013.
- Y. Su, G. Agrawal, J. Woodring, K. Myers, J. Wendelberger and J. Ahrens, "Effective and Efficient Data Sampling Using Bitmap Indices", Cluster Computing, March 2014.
- Y. Su, G. Agrawal, J. Woodring, A. Biswas and H.-W. Shen, "Supporting Correlation Analysis on Scientific Datasets in Parallel and Distributed Settings", in Proceedings of the International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC'14), June 2014, Vancouver, Canada.
- Y. Su, G. Agrawal, J. Woodring, K. Myers, J. Wendelberger and J. Ahrens. “Taming Massive Distributed Datasets: Data Sampling Using Bitmap Indices.” In Proceedings of the International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC'13), New York, NY, USA, June 2013.
- Y. Su, G. Agrawal, and J. Woodring, “Indexing and Parallel Query Processing Support for Visualizing Climate Datasets”, Proceedings of the 41st International Conference on Parallel Processing, Pittsburgh, PA, Sept. 2012.
- J. Ahrens, S. Jourdain, P. O'Leary, J. Patchett, D. H. Rogers, M. Petersen, “An Image-based Approach to Extreme Scale In Situ Visualization and Analysis”, Supercomputing 2014, New Orleans.